

RHIC Computing Facility

DRAFT

***Eric Lançon
August 23, 2016***



Outline





- Status of RCF, synergies with ATLAS Tier-1
- Performance in recent RHIC runs
- Future technological and data challenges
- Synergies with BNL Computing Initiative
- B725 infrastructure project

Status of RCF, synergies with ATLAS Tier-1

- RCF performed well during 2016 run
- Resources are ~fully utilised
- Hardware (CPU) is getting old, migration to new tape generation needed (space in HPSS)
- Synergy with ATLAS Tier-1
 - Economy of scale (operation, purchase,...)
 - Common procedures and configurations (resilience)
 - Common tools (batch system, storage, network)
 - Expertise from RCF benefit to ATLAS (and vis versa)
 - Access to Grid and cloud computing expertise





Performance in recent RHIC runs

- No issue in
 - data transfer from experiments to facility
 - Writing to tape
- give some numbers



CPU usage of the farms

- STAR farm almost continuously saturated while PHENIX farm is not
- PHENIX farm used by STAR analysis jobs
 - Optimisation of batch system (Condor) performed, STAR workflow optimisation to be done
 - STAR reconstruction cannot use PHENIX resources because STAR storage is distributed on CPU nodes
- Lessons for the future
 - Computing models (workflow management, data organisation, ...) and technological choices (storage, CPU,...) of experiments should not be too different in order to benefit from a global pool of resources

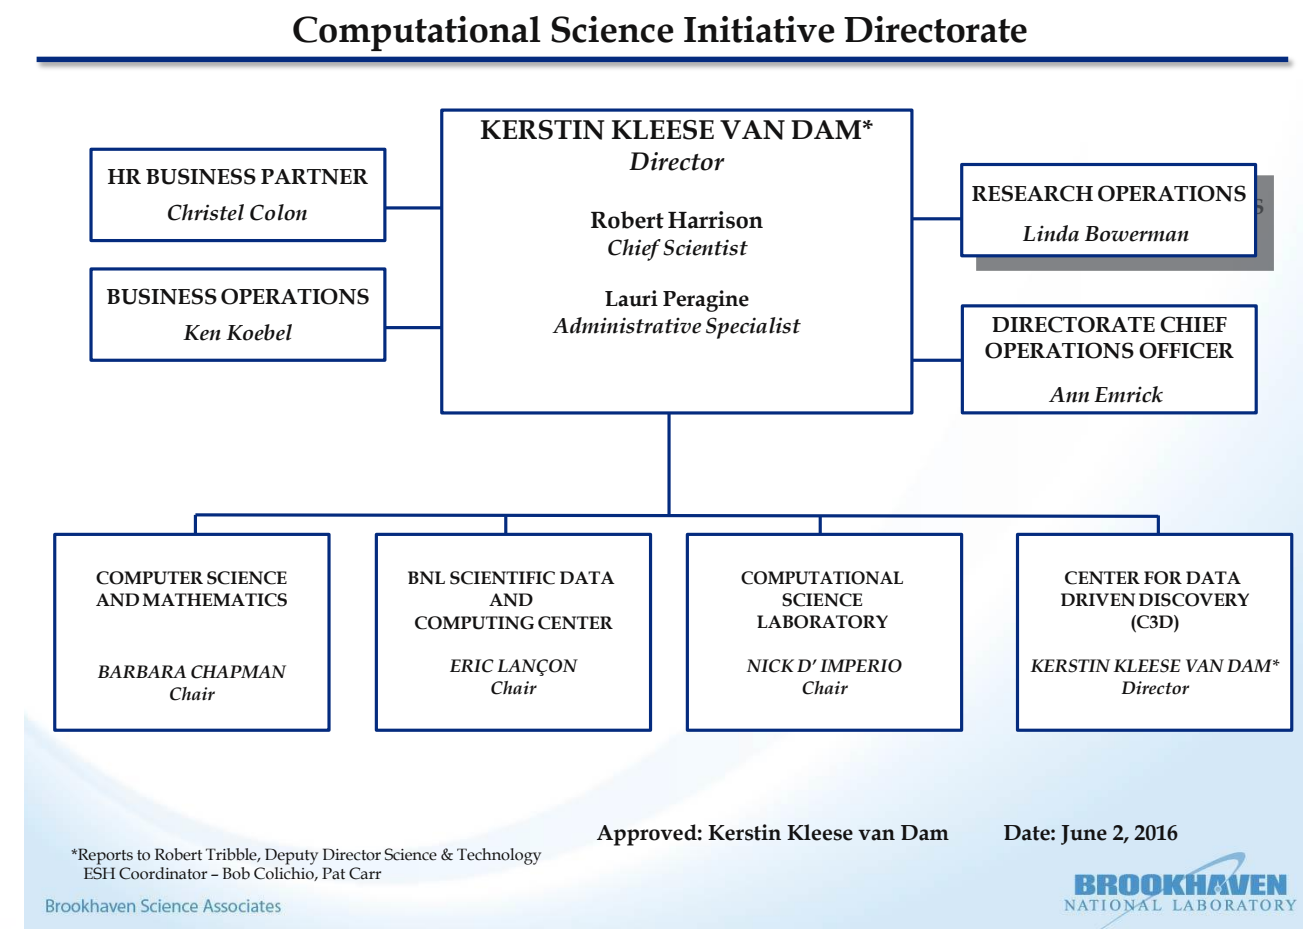


Future technological and data challenges

- Future of computing is multi-core
 - New hardware are multi-core 8, 16, 32, 64,.... with less memory per core (or at a higher price)
 - Could software of RHIC experiments be multi-core? is it worst for existing experiments?
- Object store technology, ATLAS is migrating
- RHIC hardware is getting old ~22% older than 6 years
- Tape technology
 - 2 generations behind in tape technology
 - Move to new generation (LT07) to gain in capacity / media
 - ONLY one copy of RAW on tape
 - In moving data to new media 2 copies are made
 - xxx months for migration
 - But the 2 copies will be in the same building

Computational Science Initiative : CSI

- **CSI** : *Integrating data-intensive science expertise and investments across the Laboratory to tackle "big data" challenges*
 - Leverage investments across multiple programs
 - Patterns : universities (Columbia, Cornell, New York University, Stony Brook, and Yale) and companies including IBM Research.
- Scientific Data and Computing Center : **SDCC**
- Computer...
- Computational...
- Center for....



Scientific Data and Computing Center : SDCC

- Within CSI : an integrated scientific data, computing and networking infrastructure across BNL
- RACF the RHIC and ATLAS Computing Facility is the core of SDCC :
 - A world class HEP and NP facility
 - Expertise and infrastructure necessary for leadership & innovation in data-driven discovery
- SDCC is operating
 - The Lab's new Institutional Computing system (IC)
 - A series of novel architecture research systems,
 - As well as computing and data services for other third-party clients.

Synergies with BNL Computing Initiative

- CSI is purchasing or complementing purchases in the area of HPC computing
 - Institutional cluster (Fall 2016)
 - Knight Landings (KNL) Intel farm (Fall 2016). Initiated by BNL QCD group and RIKEN, CSI doubled the capacity.
- These resources will be made available to RHIC program in opportunistic mode
- Issue : manpower to port RHIC codes?


SDCC evolution over the next months

- **Institutional Cluster (IC)**
 - Fall 2016 : 108 nodes (Intel Xeon + Nvidia K80 GPU + InfiniBand)
 - 2 x mid-2017
- **Intel Knights Landing (KNL) cluster :**
 - Fall 2016 : 144 nodes (Intel Xeon Phi + Omni-Path)
 - for RBRC (RIKEN and BNL Research Center)
 - and for CSI partners

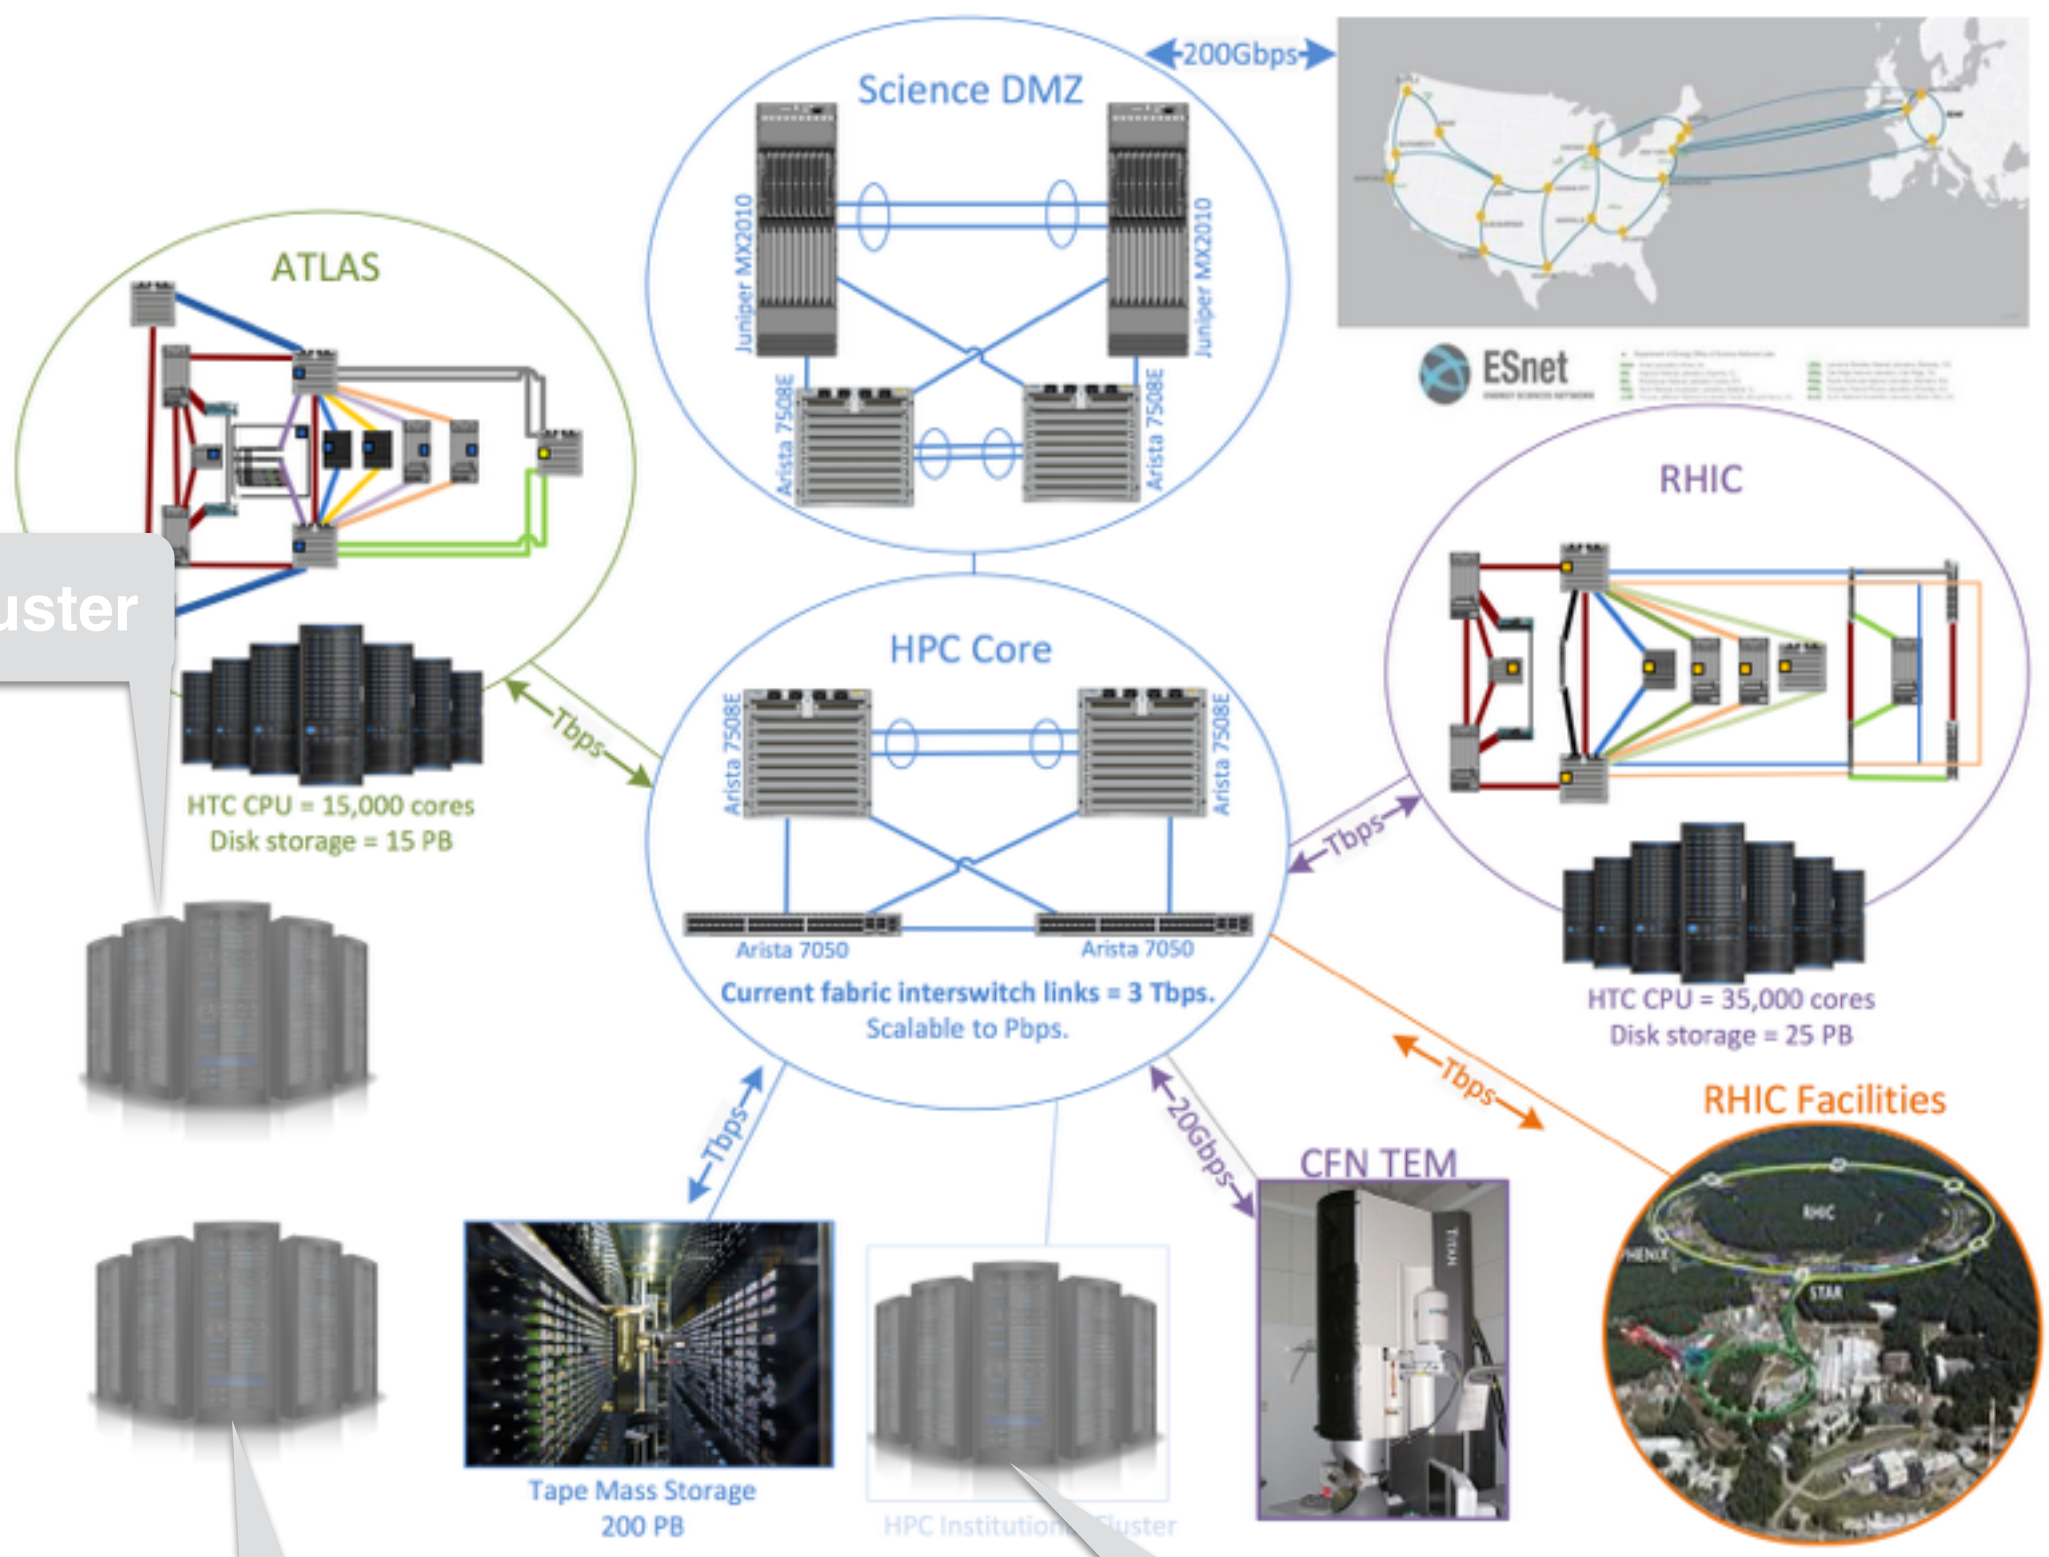
End of 2016

70k cores

HPC : 22%



KNL cluster



USQCD cluster

Institutional Cluster

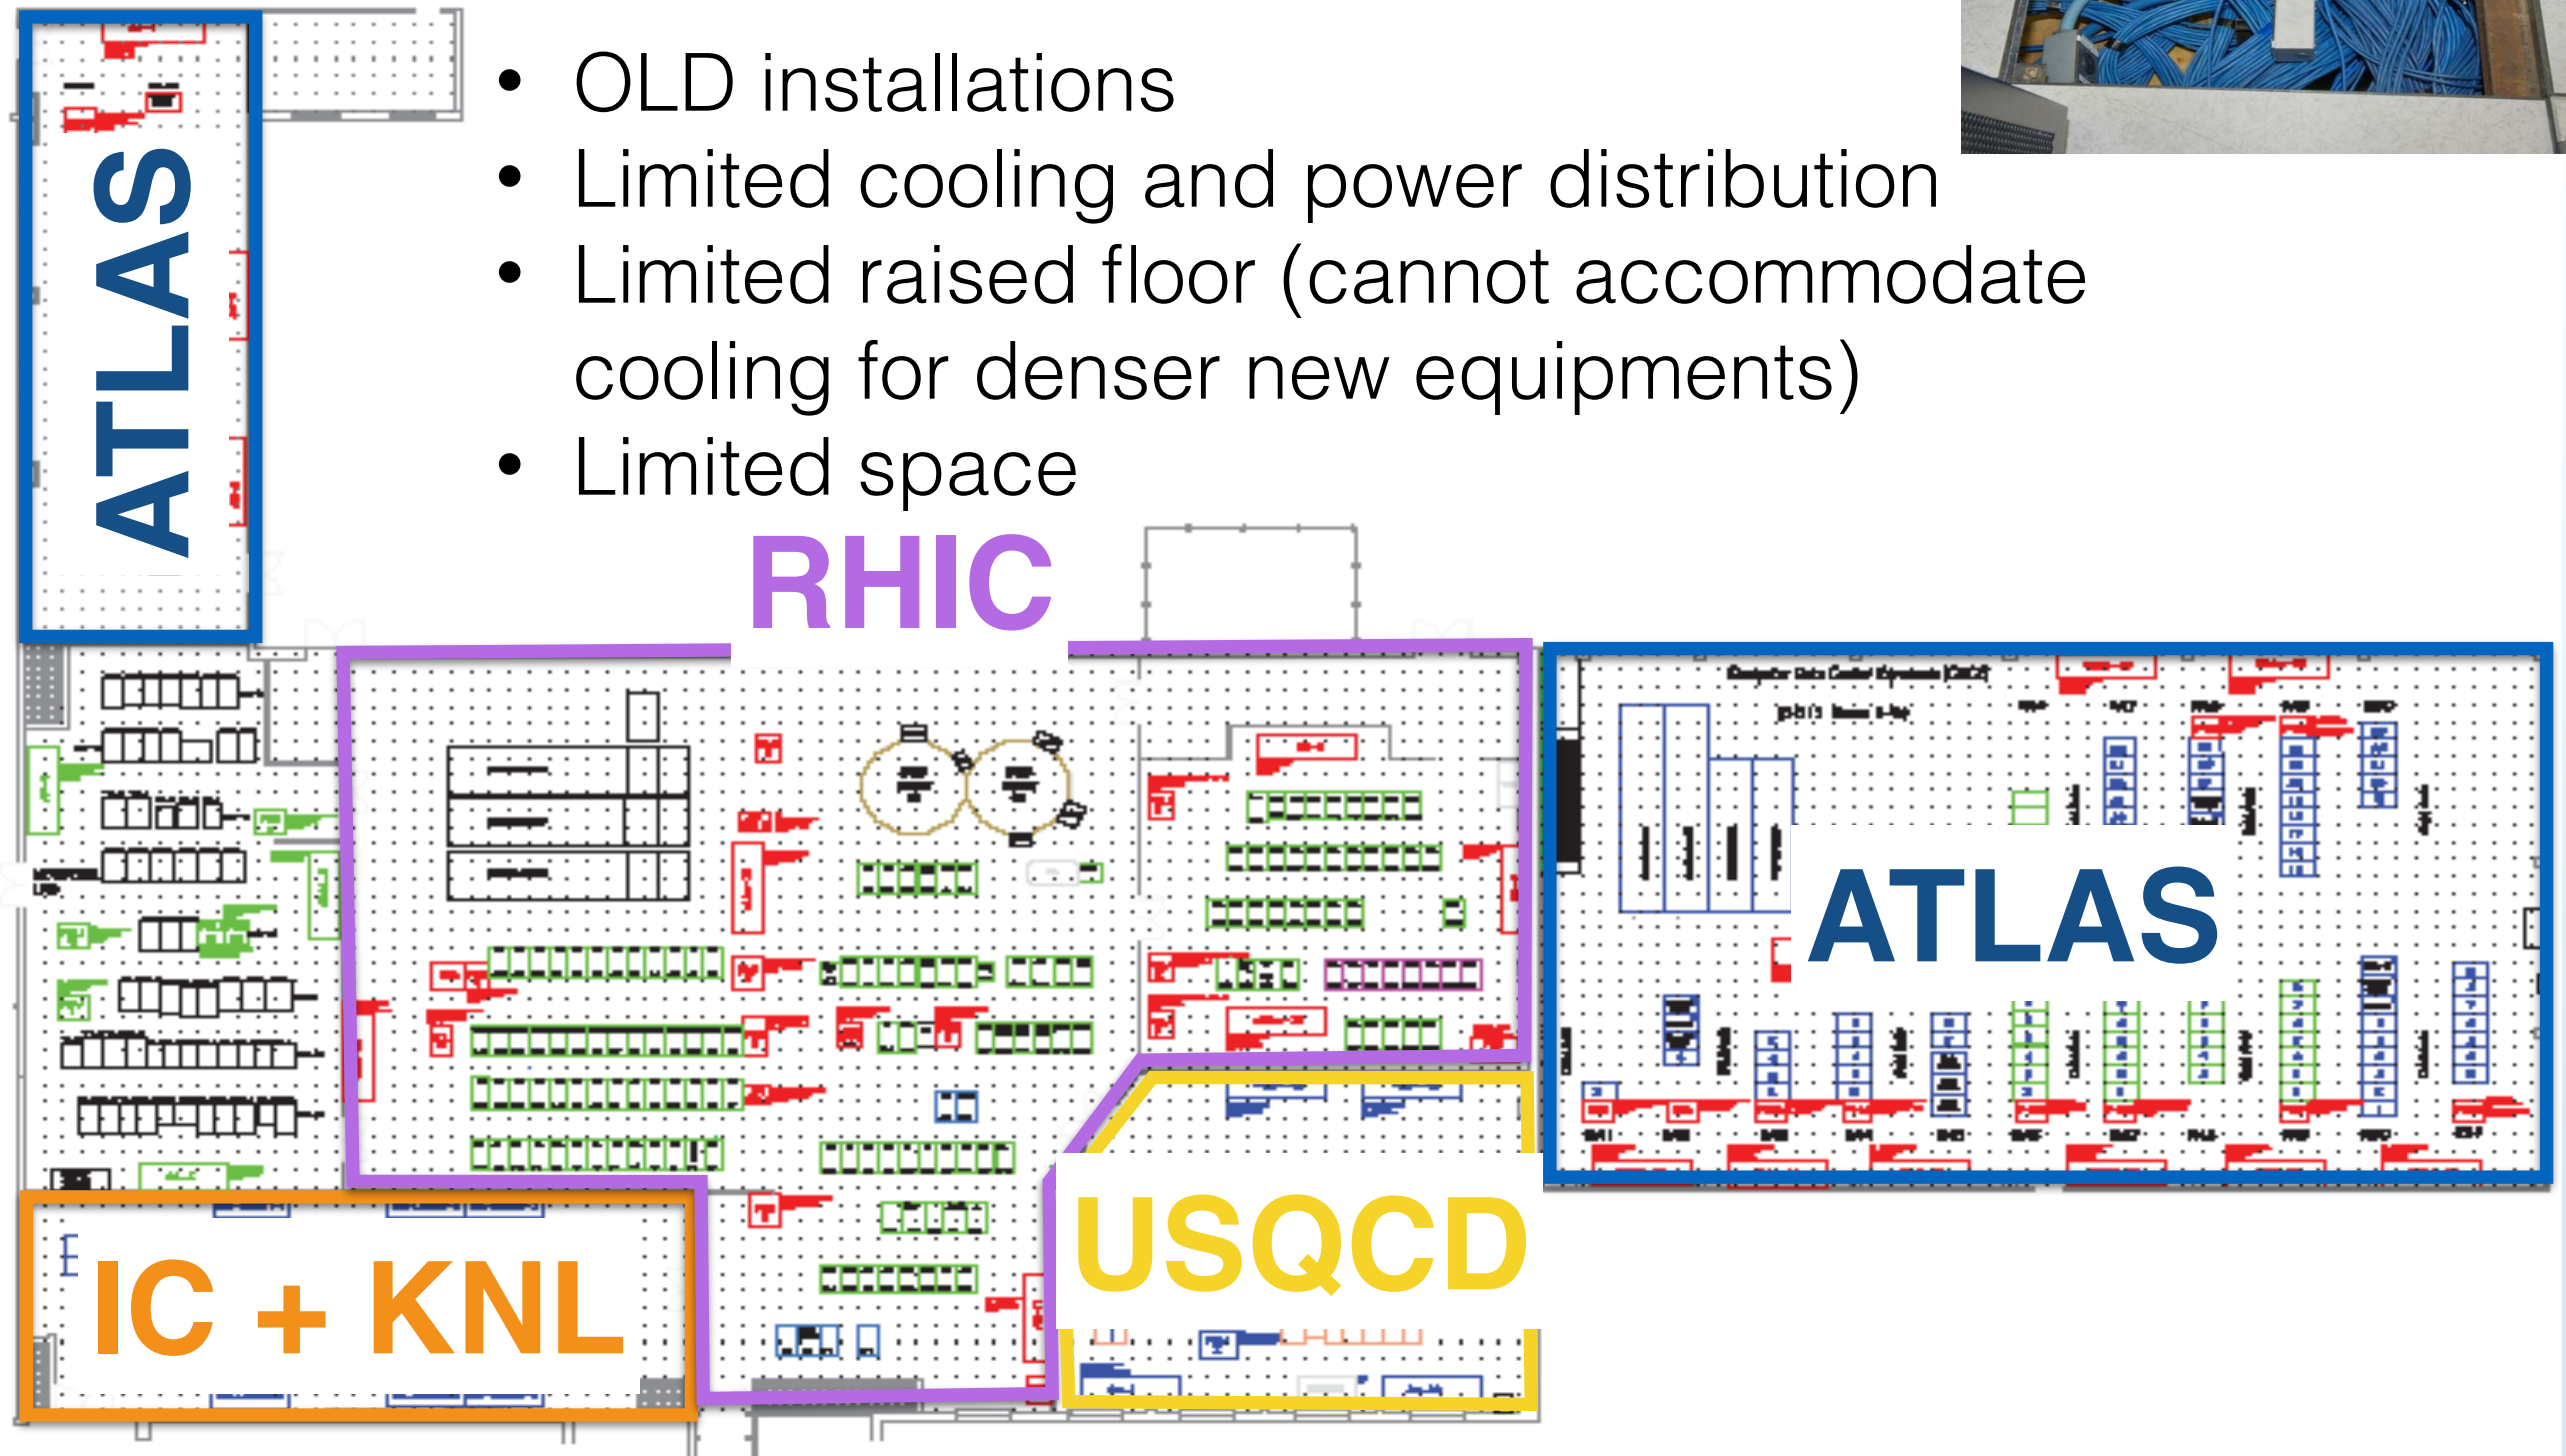
B725 infrastructure project

Computing room(s) - Issues



- OLD installations
- Limited cooling and power distribution
- Limited raised floor (cannot accommodate cooling for denser new equipments)
- Limited space

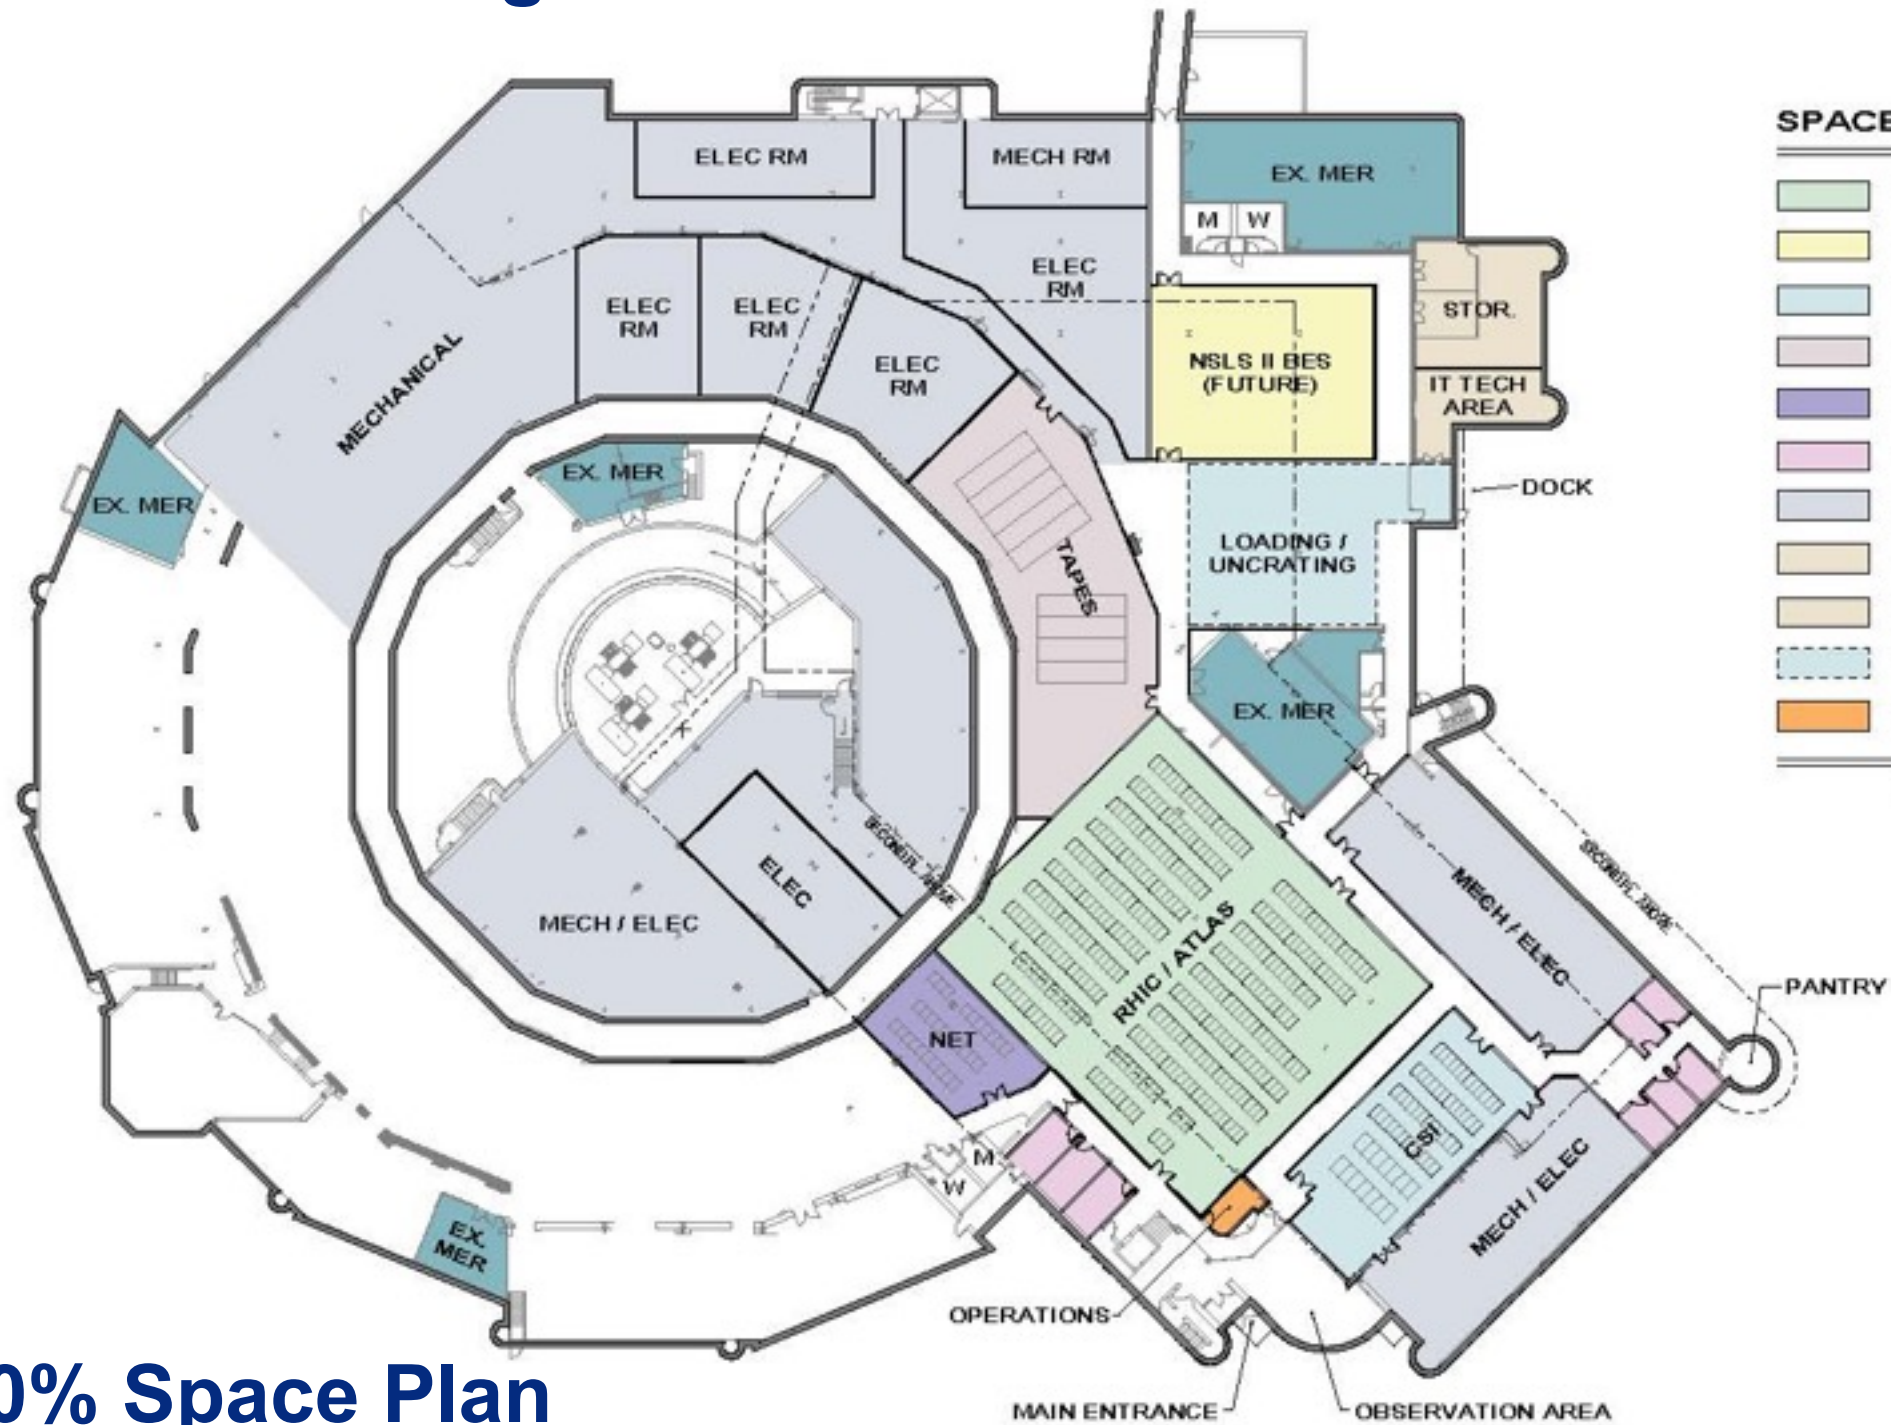
RHIC



New computing room

Core Facility Revitalisation – Conceptual Design

NSLS-I Building




SPACE LEGEND		PROGRAM NET AREA (SF)	DELIMITED AREA (SF)
	RHIC / ATLAS	10800	9271
	NSLS II BES	3400	2005
	CSI	1900	2208
	TAPES	3600	4159
	NETWORK EQUIPMENT	1800	1186
	OFFICES	1600	1186
	MECH/ELEC	26200	37504
	IT TECH AREA	800	643
	STORAGE	1200	1200
	LOADING / UNCRATING	500	2879
	OPERATIONS	100	132
TOTAL AREAS		48900	63475

CD-1 review end of August


CFR – Preliminary Schedule

Preliminary CFR Funding Analysis - 1 Yr CR 2017 (Renovation Alternative)



Longer term LHC schedule

The outline LHC schedule out to 2035 presented by Frederick Bordry to the SPC and FC June 2015 can be found [here](#)



- Delicate migration planning
- In phase with LHC Long Shutdown 2 and RHIC
 - ATLAS Tier 1 MoU : Uptime 99%

Years	Beam Species and Energies
2016	High statistics Au+Au d+Au energy scan
2017	High statistics Pol. p+p at 510 GeV
2018	⁹⁶ Zr+ ⁹⁶ Ru isobar run
2019-20	7.7-20 GeV Au+Au (BES-2)
2021	No Run ?
2022-23	200 GeV Au+Au with upgraded detectors Pol. p+p, p+Au at 200 GeV
2024---	Program TBD

CFR Design – An Incremental Approach

■ Power

- Day-one capability (2021) – 2.4 MW IT power (dedicated computing power). This is approximately double current RACF IT power.
- Provide provision for future 1.2 MW IT power increments to 6MW Max.

■ Cooling

- Day-one cooling capability to support 2.4 MW IT power
- Provide provision for future 1.2 MW IT power deployments

■ Space

- Day-one - Accommodate approximately 33% footprint expansion (Racks) within defined spaces.
- Day-one - Accommodate approximately 3,500 SF additional, unassigned space.
- Provide opportunity for future (long term) growth within the balance of the 725 facility. Both computing and offices.

Summary